

Disc-Cover Complexity Trends in Music Illustrations from Sinatra to Swift

Nicolas Fracaro
University of Padova
Italy
nicolas.fracaro@studenti.unipd.it

Stefano Cecconello
University of Padova
Italy
stefano.cecconello@unipd.it

Mauro Conti
University of Padova
Italy
Örebro University
Sweden
mauro.conti@unipd.it

Niccolò Di Marco
University of Tuscia
Italy
niccolo.dimarco@unitus.it

Alessandro Galeazzi
University of Padova
Italy
alessandro.galeazzi@unipd.it

ABSTRACT

The study of art evolution has provided valuable insights into societal change, often revealing long-term patterns of simplification and transformation. Album covers represent a distinctive yet understudied form of visual art that has both shaped and been shaped by cultural, technological, and commercial dynamics over the past century. As highly visible artifacts at the intersection of art and commerce, they offer a unique lens through which to study cultural evolution. In this work, we examine the visual complexity of album covers spanning 75 years and 11 popular musical genres. Using a diverse set of computational measures that capture multiple dimensions of visual complexity, our analysis reveals a broad shift toward minimalism across most genres, with notable exceptions that highlight the heterogeneity of aesthetic trends. At the same time, we observe growing variance over time, with many covers continuing to display high levels of abstraction and intricacy. Together, these findings position album covers as a rich, quantifiable archive of cultural history and underscore the value of computational approaches in the systematic study of the arts, bridging quantitative analysis with aesthetic and cultural inquiry.

CCS CONCEPTS

• **Applied computing** → Arts and humanities; • **Information systems** → Data mining; • **Computing methodologies** → Image processing.

KEYWORDS

Album covers, Visual complexity, Image analysis, Computational cultural analytics

ACM Reference Format:

Nicolas Fracaro, Stefano Cecconello, Mauro Conti, Niccolò Di Marco, and Alessandro Galeazzi. 2026. Disc-Cover Complexity Trends in Music Illustrations from Sinatra to Swift. In *The 41st ACM/SIGAPP Symposium on Applied*

Computing (SAC '26), March 23–27, 2026, Thessaloniki, Greece. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3748522.3779759>

1 INTRODUCTION

Art, in its various forms, has always played a central role in human civilization, shaping cultural identities and reflecting societal changes. Visual arts, including sculptures, paintings and photographs, have been extensively studied, with both qualitative and quantitative analyses revealing characteristic patterns of subject matter and visual complexity across historical periods [26, 27, 24]. Music, likewise, represents a primary cultural medium, and recent studies have highlighted a general tendency toward simplification in both musical structure and lyrics [4, 23].

Over the past century, the global expansion of the music industry and the diversification of distribution formats, from vinyl and cassettes to CDs and digital streaming, have transformed not only how music circulates but also how it is presented to audiences [3]. Album covers, initially conceived as purely promotional tools, have gradually acquired cultural and artistic significance. Much like advertisements, they both mirror and shape audience preferences, encapsulating the aesthetics of specific genres and historical contexts [10]. Beyond their artistic dimension, album covers function as cultural and communicative artifacts, comparable to other visual media. Their evolution reflects changing market strategies and audience expectations, making them a valuable case for quantitative media analysis and for understanding how cultural products visually adapt over time.

In this work, we investigate the evolution of album cover design by analyzing their visual complexity across time and musical genres. Starting from curated collections such as the Billboard Top 200, we build a large-scale dataset of 46,000 album covers spanning 75 years and multiple genres. Following previous research [26, 7, 28, 31], we apply information-theoretic and complexity-based measures to quantify their evolution, highlighting both global simplification trends and genre-specific divergences. In particular, while most genres exhibit a reduction in visual complexity, some—such as metal—show the opposite trend. Using a pre-trained machine learning-based object recognition tool, we further characterize this divergence, revealing a decreasing number of depicted elements



over time. This suggests a tendency toward growing visual minimalism in most genres, whereas others increasingly adopt abstract yet complex imagery, often featuring fused or distorted forms that hinder object recognition. Our study positions album covers as a rich, quantifiable lens on cultural and aesthetic evolution, bridging media analysis and art history, and emphasizes the value of quantitatively examining diverse art forms to achieve a deeper understanding of societal cultural shifts.

2 RELATED WORKS

Our study integrates visual-complexity analysis, machine learning-based image interpretation, and object-recognition techniques to investigate how album covers evolve across time and genres—an approach not explored in previous research. In this section, we review the key findings from each of these areas.

2.1 Visual Complexity Analysis

The notion of *visual complexity* has been extensively studied across disciplines, but remains without a universally accepted definition [5]. From a computational perspective, Kolmogorov complexity formalizes it as the length of the shortest program that can reproduce an image [14]. Since this measure is not computable in practice, compression-based metrics have been developed as practical approximations, estimating Kolmogorov complexity by exploiting the relationship between repetitiveness and compressibility [22, 19]. Alternative approaches for measuring visual complexity include frequency-domain techniques, such as Fourier or Wavelet transforms, which capture the presence of high-frequency components linked to visual detail [25], and fractal dimension measures that quantify structural irregularity [12]. More recent methods derive from information theory, including *Minimum Description Length clustering* (MDLc), which measures “meaningful complexity” by distinguishing structured signal from random noise [17], and the *Entropy-Complexity plane*, which maps images according to their degree of disorder and structural richness [26].

Beyond computational formalizations, interdisciplinary studies have applied complexity metrics to biology [1], urban evolution [2], and cultural products. In visual arts, Sigaki et al. [26] traced historical trajectories of painting styles using entropy and complexity, while in music, network-based approaches have revealed simplification trends in harmonic and melodic structures [4]. Similarly, analysis of song lyrics has documented the decrease in lexical richness and increased repetitiveness over time [23]. These findings support the idea that cultural artifacts undergo systematic simplification processes, motivating similar investigations in album cover design.

2.2 Machine Learning Approaches to Album Cover Analysis

Research specifically targeting *album covers* has primarily focused on classification and perceptual studies. Early works used low-level features such as color histograms, textures, and facial presence for genre classification [15]. Deep learning approaches, particularly convolutional neural networks (CNNs), have significantly improved classification performance by extracting high-level representations from visual data [21]. More recently, Greenfield and

Paintsil [8] leveraged balanced datasets and modern architectures such as DenseNet-201 and Vision Transformers, outperforming previous CNN-based methods. In addition, studies have highlighted correlations between simple visual features (e.g., average brightness and hue) and musical genres [6].

Beyond classification, perceptual studies have examined how cover design influences listeners’ expectations. Venkatesan et al. [30] showed that typography styles systematically affect perceived music genre: sharp-edged fonts evoke associations with energetic music (e.g., metal), while rounded fonts are linked to softer genres (e.g., pop, classical). Grebenkina et al. [7] introduced complexity metrics into this domain, demonstrating that *edge-orientation entropy* correlates positively with human aesthetic preference, including in album covers. This suggests that structural measures can approximate perceptual judgments.

Overall, prior research has either emphasized classification tasks or isolated perceptual aspects, leaving the temporal evolution of cover complexity largely unexplored. Our work addresses this gap by applying information-theoretic and machine learning-based methods to quantify how visual complexity in album covers has evolved across decades and genres.

2.3 Object Detection in Visual Art

Advanced computer vision methods have increasingly been applied to the analysis of cultural artifacts, enabling large-scale quantitative studies of visual patterns and trends [18, 13]. Recent work has applied object detection to systematic iconographic studies, enabling quantitative analysis of visual motifs across large art collections [11]. López and Flexer [16] specifically demonstrated zero-shot object detection in album covers for iconographic analysis. Our work extends this approach by using object detection to provide semantic context that enriches the interpretation of visual complexity metrics, enabling a more comprehensive understanding of visual evolution in album cover design.

3 METHODOLOGY

In this section, we describe the aggregation and preprocessing of album cover data from multiple sources, present three distinct complexity metrics to measure visual complexity, and illustrate the object detection model used to analyze semantic content.

3.1 Dataset


Data Sources and Aggregation. The dataset for this study was constructed by aggregating and standardizing data from three distinct sources. The initial foundation was built upon two existing academic datasets: the MuMu dataset¹, which provided both single- and multi-labeled genre annotations, and the MSD-I dataset², which contained single-labeled albums. To supplement these sources, particularly with more recent releases, we compiled a third dataset by scraping the Billboard charts³ from January 2002 to April 2025. This was accomplished using the `billboard.py` library to programmatically access twenty chart categories⁴. For all sources, the

¹<https://www.upf.edu/web/mtg/mumu>


²<https://zenodo.org/records/1240485>

³<https://www.billboard.com/charts/>

⁴The scraped Billboard categories include: comedy, world, new-age, reggae, christian, tropical, traditional-jazz, rock, regional-mexican, latin-rhythm, latin-pop, independent,



(a) Yearly counts of unique albums, colored by retrieval source to highlight the contribution of each dataset.



(b) Yearly genre distribution of the final dataset. Each bar represents the total number of album-genre associations for a given year, with colored segments indicating the breakdown by genre.

Figure 1: Dataset breakdown by source and genre

MusicBrainz database⁵ was used as the canonical reference to standardize album metadata, ensuring that for each entry, we retrieved the cover art corresponding to its first official release.

Deduplication and Cleaning. A multi-stage deduplication and cleaning process was implemented to ensure the uniqueness of each album. Initially, duplicates between the MuMu and MSD-I datasets were removed by cross-referencing their respective Million Song Dataset and MusicBrainz identifiers. Subsequently, albums from the Billboard source were deduplicated against the existing set using artist and title pairs. The final deduplication step was performed once all albums were mapped to a MusicBrainz ID for their first release, using this identifier as the primary key. For the cleaning phase, albums with broken image URLs or corrupted files

compilation, classical-crossover, classical, cast, americana-folk, billboard-200, vinyl, and tastemaker.

⁵<https://musicbrainz.org/>

that could not be processed by our complexity algorithms were also removed. During this process, 5,096 albums from the Billboard dataset were excluded due to faulty or missing URLs that prevented image download, while 20 images were identified as corrupted and could not be processed by the complexity estimation algorithms.

Genre Standardization. The raw genre labels aggregated from the different sources were highly heterogeneous, comprising over 1,000 unique strings that included specific sub-genres, non-genre tags (e.g., "acoustic", "1990"), and other metadata not relevant for our study. To create a consistent and meaningful set of categories for analysis, we first defined a target set of 11 primary genres⁶, guided by the taxonomy provided by Musicmap⁷. We then employed the Gemini 2.5 Pro Large Language Model via its public web-based chat interface to map each of the 1,000+ raw labels to one or more of these 11 target genres, as well as to identify and discard non-genre-related tags. This approach was chosen because the number of unique labels to be mapped was relatively small, making it feasible to complete within the interface's free usage limits.

LLM-based Data Imputation. A subset of the albums, primarily from the Billboard source, lacked genre information entirely. To address this, we performed automated genre prediction for the 2,510 unlabeled albums programmatically via the Google AI Python API. We employed the 'gemini-2.0-flash' model for this task, as it was suited for the large-scale batch processing required for our analysis and was the freely accessible model via the API at the time of the research. We designed a structured prompt that constrained the model to classify albums using only our predefined 11 supergenres and return results in a standardized CSV format with confidence indicators⁸. To validate this approach, we tested the model on a control set of 500 albums with known labels, instructing it to predict a single genre from our set of 11 and to provide a boolean confidence value. The model reported high confidence for 450 predictions, of which 87% were correctly contained within the album's true genre labels. For the 50 low-confidence predictions, this accuracy dropped to 62%. Based on these results, we proceeded with the imputation but only accepted high-confidence predictions. Finally, 911 albums in the dataset were missing a release date. We explored using the 'gemini-2.0-flash' model to predict the year of release. However, a validation test on 500 albums with known dates yielded unsatisfactory results. The model returned high-confidence predictions for only 289 cases. For this subset, the accuracy for predicting the exact year was 64.4%, rising to 80.3% for a tolerance of ± 5 years. Given

⁶The genres include: Pop, Rock, Speciality, Jazz & Blues, Electronic, World Music, Country & Folk, Metal, R&B, Hip Hop, Classical

⁷Musicmap. The Genealogy and History of Popular Music Genres. <https://musicmap.info>. Accessed: July 2025.

⁸Full prompt: "You are a music expert tasked with classifying album genres. I will provide you with a list of albums, and you need to classify each one using ONLY the following 11 supergenres: Classical, Country & Folk, Electronic, Hip Hop, Jazz & Blues, Metal, Pop, R&B, Rock, Speciality, World music IMPORTANT INSTRUCTIONS: 1. Return ONLY a CSV format with the following columns: album_group_mbid, title, release_group_date, genres, sure 2. The genres column should contain a Python list format: ['Genre'], only 1 single genre per album is allowed 3. Use ONLY the 11 supergenres listed above - no other genres are allowed 4. If you're unsure about a genre, pick the most likely one from the 11 options but set "sure" to false 5. Do not include any explanations, headers, or additional text - ONLY the CSV data 6. Do not include markdown formatting or code blocks 7. CRITICAL: Use proper CSV escaping - put double quotes around any field that contains commas, quotes, or special characters 8. Example format: "album_id","Title, with comma","2023",["Pop"],"true" Here are the albums to classify:

the low confidence rate and insufficient accuracy, we deemed this approach unreliable for our time-series analysis. Consequently, we opted to remove these 911 albums from the dataset to maintain the temporal integrity of our analysis.

Final Dataset Characteristics. After the processes of cleaning, deduplication, and standardization, the final dataset consists of 46,399 unique, multi-labeled albums, each with a corresponding cover image. The dataset composition is presented in Figure 1a, which reports the number of unique albums per year by source. On the other hand, Figure 1b illustrates the genre distribution over time, with albums associated with multiple genres counted more than once. It is important to note the presence of two main biases: a genre bias, with a predominance of Pop and Rock albums, and a temporal bias, with the majority of albums concentrated between 1990 and 2010. Furthermore, as a multi-labeled dataset, a single album can be associated with, and thus contribute to, more than one genre category. A specific feature of the temporal distribution is the decrease in the number of albums between 1999 and 2003. This trend was already present in the publicly available MuMu and MSD-I datasets, and became slightly more pronounced when extending coverage with Billboard data starting from 2002. Despite these biases, our final dataset provides a reasonable representation of the musical landscape, encompassing the most popular genres over a span of 75 years.

3.2 Complexity Metrics

As no single metric can capture all aspects of complexity, we selected distinct computational methods that provide a comprehensive overview, complemented by an object detection analysis to add a semantic dimension.

Entropy-Complexity Plane. To disentangle visual disorder from structural richness, we utilized the Entropy-Complexity plane methodology developed by Sigaki et al. [26]. This approach was chosen for its ability to characterize an image using two complementary measures, both bounded between 0 and 1: Permutation Entropy (H), which quantifies the degree of randomness in local pixel patterns, and Statistical Complexity (C), which measures the presence of elaborate spatial structures. An image is thus represented as a point in a 2D plane, allowing for a more nuanced analysis than a single complexity score. For example, an image of random noise would exhibit high entropy (approaching 1) due to its unpredictable pixel patterns, but low statistical complexity (near 0) because it lacks coherent spatial structures. A solid white image would share this low statistical complexity due to its lack of spatial structure, while showing near-zero entropy because of its complete predictability. Following state of the art techniques, each image was converted to its grayscale representation as a pre-processing step. The calculation is then based on the distribution of local ordinal patterns (2x2 pixel windows in our implementation), as described in the original paper. To illustrate the meaning of the two measures in some real-case examples, Figure 2 presents the covers of four famous musical albums, positioned according to their values of statistical complexity and permutation entropy.

Meaningful Complexity (MDLc). To measure the “meaningful” complexity of an image, we implemented the method proposed by

Mahon and Lukasiewicz [17]. MDLc values range from 0 (indicating minimal meaningful structure) to theoretically unbounded positive values, where higher scores represent more complex images. We selected this metric for its specific design to distinguish structured, significant information from random noise. Unlike traditional metrics that might assign high complexity to a textureless, noisy image, MDLc uses a hierarchical clustering approach based on the Minimum Description Length (MDL) principle to identify and quantify coherent structures at multiple scales. For our implementation, we used the authors’ publicly available code⁹. All images were resized to a standard resolution of 224x224 pixels to ensure consistency. To make the computation feasible on our large dataset, we reduced the maximum number of clusters (K_{max}) to test during the MDL search from 8 to 5, a modification suggested by the original authors to improve efficiency without significantly impacting accuracy. All other hyperparameters were kept at their default values, and no training or fine-tuning procedures have been performed.

Compression-based Complexity (ZIPc). As a simple and objective baseline, we included a compression-based metric, ZIPc. This approach provides a practical approximation of Kolmogorov complexity, where the complexity of an image is estimated by its incompressibility. We chose this metric because it offers a robust measure of statistical redundancy that is straightforward to compute and interpret. The ZIPc score was calculated as the ratio of the compressed file size to the original raw bitmap size. The metric typically produces values between 0 and 1, where 0 represents perfect compressibility (minimal complexity) and values approaching 1 indicate high incompressibility (maximum complexity). In rare cases, values may exceed 1 when compression overhead surpasses any achieved compression gains. To ensure a standardized process, each image was first converted to an uncompressed RGB bitmap. The bitmap data was then compressed using the DEFLATE algorithm (via the ZIP file format) with the compression level set to its maximum value (9).

Semantic Complexity via Object Detection. To complement the perceptual complexity metrics, we introduced a measure of semantic complexity based on the number of recognizable objects in an image. We employed the YOLOv8 model [9], specifically the lightweight yolov8n.pt version pre-trained on the COCO dataset. This model was chosen for its high efficiency and accuracy in identifying a predefined set of 80 common object classes. For each album cover, we extracted both the total count of detected objects and their corresponding object classes using a confidence threshold of 0.25. Object counts range from 0 (no detected objects) to theoretically unbounded positive integers. This approach provides both a quantitative measure of the “busyness” of album covers in terms of recognizable content and qualitative insights into the types of objects commonly featured across different genres.

4 RESULTS AND DISCUSSION

This section presents and discusses the results of the analysis of temporal evolution and genre differences in album cover complexity using three quantitative metrics, supplemented by automated object

⁹https://github.com/Lou1sM/meaningful_image_complexity?tab=readme-ov-file

detection to identify the semantic drivers underlying the observed patterns.

4.1 Entropy-Complexity Plane Metric

We begin our analysis by considering the entropy–complexity plane. As discussed in Section 3, this representation provides a nuanced perspective on the visual characteristics of album covers, revealing distinct patterns across genres and over time. To build intuition, Figure 2 presents four representative examples that illustrate the visual properties associated with different regions of the plane. The two columns correspond to low (left) and high (right) entropy, while the two rows represent low (bottom) and high (top) complexity. Low-entropy covers are generally characterized by uniform backgrounds and a small number of orderly visual elements, in sharp contrast to the random-like noise associated with high entropy. Within this group, the Whiskey Myers album (bottom left) is visually simple, consisting only of a text element, whereas the Pink Floyd album (top left) demonstrates higher complexity through highly regular patterns, such as the triangular shape and colored stripes. In contrast, high-entropy covers are marked by more chaotic, blurred contours and shaded elements, which increase their resemblance to random noise. In this region, background structure plays a decisive role in differentiating levels of complexity. For example, the Bob Dylan album (top right) exhibits higher statistical complexity due to structured patterns such as the brick wall and scarf, while the Red Hot Chili Peppers album (bottom right) lacks comparable regularity and therefore presents lower complexity despite its high entropy.

Figure 3a shows the Entropy-Complexity plane grouped by genres, where each point represents the mean entropy and complexity values for that genre. Interestingly, the figure reveals a dense central cluster comprising a majority of the music genres. The overlapping standard error bars confirm the statistical significance of this grouping, which points to a widely adopted stylistic norm of moderate permutation entropy and statistical complexity. Beyond this central cluster, three genres emerge as distinct outliers. Metal and Hip Hop are characterized by higher entropy and lower statistical complexity, a profile that may reflect more visually disordered aesthetics. Conversely, the Electronic genre occupies an opposing region of lower entropy and higher statistical complexity, which suggests more structured, minimalist designs.

We repeat the analysis grouping albums over time. The results are shown in Figure 3b, where time is divided into dynamic periods containing approximately 3000 albums each. This threshold balances data sufficiency with temporal granularity, ensuring each period contains adequate data for stable comparisons while preserving the ability to detect temporal trends. The results reveal a clear, non-random evolutionary trajectory that can be grouped into four macro-periods. Initially, from the 1950s to 1996, a rise in permutation entropy and a fall in statistical complexity suggest a trend toward more disordered artwork. This pattern shifted between 1997 and 2003, as lower entropy and higher statistical complexity indicate a shift toward more structured designs. A period of relative stability followed from 2004 to 2009. Finally, the most recent period (2010–2025) is marked by a sharp decrease of approximately 12% in statistical complexity at stable entropy levels. Since entropy remains largely unchanged, indicating consistent visual disorder,

the observed decline in statistical complexity reflects a shift toward compositions with reduced structural elaboration. This trajectory points to a move toward more streamlined designs, with fewer intricate spatial patterns in the modern era.


Several factors may be at the core of the observed trend, encompassing both cultural shifts and technological developments. Culturally, the rising popularity of metal and hip hop during the 1980s and 1990s may have influenced the visual style of album covers in other genres, contributing to increased permutation entropy alongside reduced statistical complexity during this period. From a technological perspective, the emergence of music-oriented digital platforms, such as iTunes (launched globally in 2004), may have driven the decrease in statistical complexity from 2004 onward. Indeed, these platforms typically displayed album artwork as small thumbnails—by 2010, limited to 600×600 pixels at 72 dpi [20]—potentially discouraging the inclusion of highly detailed designs.




Figure 2: Album cover examples with different complexity and entropy scores. (top left) Pink Floyd, *The Dark Side of the Moon*, 1973. (top right) Bob Dylan, *Blonde on Blonde*, 1966. (c) Whiskey Myers, *Whiskey Myers*, 2019. (d) By the Way, Red Hot Chili Peppers, 2002.

4.2 MDLc and ZIPc Complexity Metrics

Analysis of the MDLc, which measures meaningful information by distinguishing coherent patterns from random noise, and ZIPc, which uses compression ratios as an estimate of image complexity, corroborates the findings from the entropy-complexity plane, confirming the notable similarity in visual complexity across most genres, as shown in Figure 4. These metrics show that, while most genres maintain proximate complexity scores throughout the observed timeline, the three previously identified outliers are consistent across these measures: Metal and Hip Hop consistently exhibit the highest average complexity scores, while the Electronic genre



(a) Mean Permutation Entropy and Statistical Complexity by genre.




(b) Evolutionary trajectory of mean complexity values over time. Periods were defined of variable duration to ensure a minimum of 3000 albums per period (except the final period, which contains 2,269)

Figure 3: Characterization of album covers in the Entropy-Complexity plane. Panel (a) shows the average positioning for each music genre, highlighting a central cluster and key outliers. Panel (b) illustrates the aggregate temporal evolution across all genres.


consistently registers among the lowest. These results reinforce the characterization of Metal and Hip Hop's aesthetics as visually dense, and Electronic's as minimalist.

The temporal evolution of the MDLc and ZIPc scores also aligns with the four macro-periods identified previously. A general increase in complexity is observed until the mid-1990s, followed by a reversal and a period of stability. Critically, both metrics confirm the significant, widespread decrease in average complexity during the most recent period (2010-2025), lending further support to the hypothesis of a shift towards less elaborate cover art. This finding confirms that the general trend towards simplification, previously

documented for lyrical content [23, 29] and melodic structure [4], also extends to the visual dimension of music. This convergence towards simplicity across these domains suggests a shared response to the modern media landscape, where a plausible driver is the shift to digital platforms. In a context where immediate engagement is critical, simpler, bolder designs may be favored for their ability to capture attention quickly, particularly on small screens or in recommendation-driven listening environments, just as more repetitive lyrics and less complex melodies serve similar functions in the auditory domain.



(a) Average MDLc score over time.



(b) Average ZIPc score over time.

Figure 4: Temporal evolution of average visual complexity across genres, measured by (a) MDLc and (b) ZIPc. Periods were defined of variable duration to ensure a minimum of 3000 albums per period (except the final period, which contains 2,269). A genre is included in a period only if at least 50 albums of that genre are present.

However, this trend toward simplification is nuanced by a simultaneous increase in stylistic diversity. While the median complexity has decreased, the dispersion of complexity scores has progressively widened over time, within each period comprising at least 3000 albums, as illustrated in Figure 5. This is most pronounced in the

last 15 years, during which the interquartile range has increased by approximately 11% for both the MDLc and ZIPc metrics. This suggests that the modern media landscape has not led to visual homogenization. Instead, it appears to foster a polarization of aesthetics, where a dominant trend toward minimalist, high-impact designs coexists with the continued exploration of highly complex, niche visual styles. It is plausible that the digital landscape itself creates favorable conditions for this diversity; the democratization of advanced creative software facilitates the production of intricate visuals, while the lowered barrier to entry on digital platforms may reduce the commercial pressures of traditional record labels, potentially creating more space for such artistic experimentation. Interestingly, most of the outliers appear on the lower end of the distribution. One possible interpretation is that this asymmetry may reflect an intrinsic characteristic of album cover design: while there might be aesthetic constraints on how far visual complexity can be pushed, it may be comparatively easier for designs to deviate toward the minimalist extreme.

4.3 Interpreting Complexity through Object Detection

Previously, we used complexity metrics to capture broad patterns in album covers, without detailing the shifts behind them. Here, we address this by analyzing semantic content using the YOLOv8 model to extract object counts and classes. By applying object detection on album cover, we analyze the trend on the presence and classes over time and across genres.

As shown in Figure 6a, the “person” class is the most frequent semantic element across all genres. This widespread use of portraiture, which could serve as a direct marketing tool by placing the artist at the forefront, may help explain the moderate, tightly-grouped complexity scores that form the stylistic norm, potentially reflecting constraints on visual language diversity. Using the same dynamic periods of at least 3000 albums each, the temporal evolution of object counts (Figure 6b) further clarifies the nature of the outlier genres. The case of Metal is particularly insightful: despite its high complexity, our object detection analysis reveals it has the lowest average number of detected objects. This apparent paradox may indicate that Metal’s complexity derives from sources other than recognizable objects, potentially abstract compositions, intricate logos, and detailed textures. In contrast, the low object count for the Electronic genre aligns with its low complexity scores, corroborating the interpretation of a minimalist aesthetic based on few recognizable elements.

5 CONCLUSIONS

We study the evolution of album covers as art and as cultural artifact, emphasizing how they have changed in response to digital technologies. Consistent with previous studies [4, 23, 26], our analysis, based on multiple measures of visual complexity and object recognition methods, reveals an overall simplification of album covers over time, with notable exceptions in genres such as Metal and Hip Hop. A closer look suggests that more experimental genres tend to increase complexity through abstraction, whereas other genres rely on artists’ images, likely as a promotional tool, resulting in lower visual complexity. Future work could examine the role





Figure 5: Evolution of the distribution of complexity scores over time for (a) the MDLc metric and (b) the ZIPc metric. While the median complexity (red dashed line) trends downward in recent years, the interquartile range (height of the boxes) increases, indicating growing stylistic diversity.


of color in the Entropy–Complexity plane, adopt more advanced models to capture abstract symbolism, and compare the trajectories of different cover editions with global trends.

More broadly, investigating the visual complexity of album covers contributes to understanding how cultural products circulate and acquire meaning. Since platforms such as social networks rely heavily on visual cues to capture attention, trends in design and simplification can influence not only aesthetic choices but also the visibility and reception of cultural artifacts. By linking visual complexity to broader processes of cultural communication, this type of study provides valuable insights for media analysis, highlighting how artistic and promotional strategies intersect with the dynamics of digital platforms.

Acknowledgments A.G. gratefully acknowledges the financial support provided by CY4GATE. A.G. and S.C. acknowledge Prof. Daniela Tonon for making this collaboration possible.



(a) Proportional distribution of detected object classes per genre, including albums with no detected objects and the top 10 most frequently detected object classes.



(b) Average number of detected objects per album over time.

Figure 6: Analysis of semantic content via object detection across music genres. Panel (a) shows the proportional distribution of detected object classes per genre, including albums with no detected objects. The ‘person’ class predominates across all genres, while Electronic and Metal show higher proportions of albums with no detected objects. Panel (b) illustrates the temporal evolution of average object counts per album. Metal consistently exhibits the lowest object count despite high visual complexity, while Electronic maintains low counts consistent with its minimalist aesthetic.

REFERENCES

- [1] Christoph Adami. 2002. What is complexity? *BioEssays*, 24, 12, 1085–1094.
- [2] Michael Batty. 2016. Complexity in city systems: understanding, evolution, and design. In *A planner's encounter with complexity*. Routledge, 99–122.
- [3] Johannes Brusila, Martin Cloonan, and Kim Ramstedt. 2021. Music, digitalization, and democracy. *Popular Music and Society*, 45, 1, (Oct. 2021), 1–12. doi: 10.1080/03007766.2021.1984018.
- [4] Niccolò Di Marco, Edoardo Loru, Alessandro Galeazzi, Matteo Cinelli, and Walter Quattrociochi. 2025. Decoding musical evolution through network science. *arXiv preprint arXiv:2501.07557*. arXiv: 2501.07557 [cs. SD] (<http://cs.sd/>).
- [5] Don C Donderi. 2006. Visual complexity: a review. *Psychological bulletin*, 132, 1, 73.
- [6] Aleksandra Dorochowicz and Bożena Kostek. 2019. Relationship between album cover design and music genres. In *2019 Signal Processing: Algorithms, Architectures, Arrangements, and Applications (SPA)*. IEEE, 93–98.
- [7] Maria Grebenkina, Anselm Brachmann, Marco Bertamini, Ali Kaduham, and Christoph Redies. 2018. Edge-orientation entropy predicts preference for diverse types of man-made images. *Frontiers in Neuroscience*, 12, 678.
- [8] Kaelen Greenfield and Ebenezer Paintsil. 2024. Album covers deserve some attention.
- [9] Glenn Jocher, Ayush Chaurasia, Jing Qiu, and et al. 2023. YOLOv8. <https://github.com/ultralytics/ultralytics>. Accessed 14th July 2025. (2023).
- [10] Steve Jones and Martin Sorger. 1999. Covering music: a brief history and analysis of album cover design. *Journal of popular Music studies*, 11, 1, 68–102.
- [11] Anne Marie Kanstrup et al. 2024. Algorithmic ways of seeing: using object detection to facilitate art exploration. In *CHI '24*.
- [12] Nina Siu-Ngan Lam, Hong-lie Qiu, Dale A Quattrocchi, and Charles W Emerson. 2002. An evaluation of fractal methods for characterizing image complexity. *Cartography and Geographic Information Science*, 29, 1, 25–35.
- [13] Sabine Lang and Björn Ommer. 2021. Transforming information into knowledge: how computational methods reshape art history. *Digital Humanities Quarterly*, 15, 3.
- [14] Ming Li, Paul Vitányi, et al. 2008. *An introduction to Kolmogorov complexity and its applications*. Vol. 3. Springer.
- [15] Janis Libeks and Douglas Turnbull. 2011. You can judge an artist by an album cover: using images for music annotation. *IEEE MultiMedia*, 18, 4, 30–37.
- [16] Selene López and Arthur Flexer. 2025. Zero-shot open set object detection in music album covers. In (July 2025). doi: 10.5281/zenodo.15839738.
- [17] Louis Mahon and Thomas Lukasiewicz. 2024. Minimum description length clustering to measure meaningful image complexity. *Pattern Recognition*, 145, 109889. doi: <https://doi.org/10.1016/j.patcog.2023.109889>.
- [18] Lev Manovich. 2020. *Cultural Analytics*. MIT Press.
- [19] Manuela M. Marin and Helmut Leder. 2013. Examining complexity across domains: relating subjective and objective measures of affective environmental scenes, paintings and music. *PLoS ONE*, 8, 8, e72412. doi: 10.1371/journal.pone.0072412.
- [20] Chris McVeigh. 2010. Creating original cover art for itunes. Retrieved Sept. 25, 2025 from <https://www.macworld.com/article/206005/itunesart.html>.
- [21] Sergio Oramas, Oriol Nieto, Mohamed Sordo, and Xavier Serra. 2018. Multimodal deep learning for music genre classification. *Transactions of the International Society for Music Information Retrieval*, 1, 1, 4–21. doi: 10.5334/tismir.10.
- [22] Letizia Palumbo, Ruth Ogden, Alexis D.J. Makin, and Marco Bertamini. 2015. Computerized measures of visual complexity. *Behavior Research Methods*, 47, 3, 895–909. doi: 10.3758/s13428-014-0533-y.
- [23] Emilia Parada-Cabaleiro, Maximilian Mayerl, Stefan Brandl, Marcin Skowron, Markus Schedl, Elisabeth Lex, and Eva Zangerle. 2024. Song Lyrics Have Become Simpler and More Repetitive Over the Last Five Decades. *Scientific Reports*, 14, 1, 5531. Erratum in: *Sci Rep*. 2024 May 22;14(1):11712. doi: 10.1038/s41598-024-55742-x.
- [24] POURIA PEDRAM and G. R. JAFARI. 2008. Mona lisa: the stochastic view and fractality in color space. *International Journal of Modern Physics C*, 19, 06, (June 2008), 855–866. doi: 10.1142/s0129183108012558.
- [25] Christoph Redies, Jens Hasenstein, and Joachim Denzler. 2007. Fractal-like image statistics in visual art: similarity to natural scenes. *Spatial vision*, 21.
- [26] Igor Y. D. Sigaki, Matjaž Perc, and Haroldo V. Ribeiro. 2018. History of art paintings through the lens of entropy and complexity. *Proceedings of the National Academy of Sciences of the United States of America*, 115, 37, E8585–E8594. doi: 10.1073/pnas.1800083115.
- [27] Richard P. Taylor, Adam P. Micolich, and David Jonas. 1999. Fractal analysis of pollock's drip paintings. *Nature*, 399, 6735, (June 1999), 422–422. doi: 10.1038/20833.
- [28] Carlo M Valensise, Alessandra Serra, Alessandro Galeazzi, Gabriele Etta, Matteo Cinelli, and Walter Quattrociochi. 2021. Entropy and complexity unveil the landscape of memes evolution. *Scientific Reports*, 11, 1, 20022.
- [29] Michael E. W. Varnum, Jaimie A. Krems, Cavan Morris, Alexandra Wormley, and Igor Grossmann. 2021. Why are song lyrics becoming simpler? a time series analysis of lyrical complexity in six decades of american popular music. *PLoS ONE*, 16, 1, e0244576. doi: 10.1371/journal.pone.0244576.
- [30] Tara Venkatesan, Qian Janice Wang, and Charles Spence. 2022. Does the typeface on album cover influence expectations and perception of music? *Psychology of Aesthetics, Creativity, and the Arts*, 16, 3, 487.
- [31] Haoyu Wang, Changqing Song, and Peichao Gao. 2024. Complexity and entropy of natural patterns. *PNAS nexus*, 3, 10, pgae417.